

High Dimensional Forecasting via Interpretable Vector Autoregression*

William B. Nicholson[†], Jacob Bien[‡] and David S. Matteson[§]

Cornell University

January 28, 2016

Abstract

Vector autoregression (VAR) is a fundamental tool for modeling multivariate time series. However, as the number of component series is increased, the VAR model becomes overparameterized. Several authors have addressed this issue by incorporating regularized approaches, such as the lasso in VAR estimation. Traditional approaches address overparameterization by selecting a low lag order, based on the assumption of short range dependence, assuming that a universal lag order applies to all components. Such an approach constrains the relationship between the components and impedes forecast performance. The lasso-based approaches work much better in high-dimensional situations but do not incorporate the notion of lag order selection.

We propose a new class of regularized VAR models, called hierarchical vector autoregression (HVAR), that embed the notion of lag selection into a convex regularizer. The key modeling tool is a group lasso with nested groups which guarantees that the sparsity pattern of lag coefficients honors the VAR's ordered structure. The HVAR framework offers three structures, which allow for varying levels of flexibility. A simulation study demonstrates improved performance in forecasting and lag order selection over previous approaches, and two macroeconomic applications further highlight forecasting improvements as well as HVAR's convenient, interpretable output.

1 Introduction

Vector autoregression (VAR) has emerged as the standard-bearer for macroeconomic forecasting since the seminal work of Sims (1980). VAR is also widely applied in numerous fields, including climatology, neuroscience, and signal processing. The number of VAR parameters grows quadratically with the the number of component series, and, in the words of Sims, this “profligate parameterization” becomes intractable for large systems of variables. Without further assumptions, VAR modeling is infeasible except in limited situations in which the number of components and the lag order are small.

Many approaches have been proposed for reducing the dimensionality of vector time series models, including canonical correlation analysis (Box & Tiao 1977), factor models (Peña & Box 1987), scalar component models (Tiao & Tsay 1989),

*The authors thank Gary Koop for providing his data transformation script. This research was supported by an Amazon Web Services in Education Research Grant. JB was supported by NSF DMS-1405746 and DSM was supported by NSF DMS-1455172 and a Xerox Corporation Ltd. faculty research award.

[†]PhD Candidate, Department of Statistical Science, Cornell University; email: wbn8@cornell.edu; url: <http://www.wbnicholson.com>

[‡]Assistant Professor, Department of Biological Statistics and Computational Biology and Department of Statistical Science, Cornell University, 1178 Comstock Hall, Ithaca, NY 14853; email: jbien@cornell.edu; url: <http://faculty.bscb.cornell.edu/~bien/>

[§]Assistant Professor, Department of Statistical Science and ILR School Department of Social Statistics, Cornell University, 1196 Comstock Hall, Ithaca, NY 14853; email: matteson@cornell.edu; url: <http://stat.cornell.edu/~matteson/>

independent component analysis (Back & Weigend 1997), principal component analysis (Stock & Watson 2002), generalized dynamic factor models (Forni et al. 2000), and dynamic orthogonal component models (Matteson & Tsay 2011).

Many recent approaches have instead focused on imposing sparsity in the estimated coefficient matrices through the use of convex regularizers such as the lasso (Tibshirani 1996). Most of these methods are adapted from the standard regression setting and do not specifically leverage the ordered structure inherent to the lag coefficients in a VAR. We propose a new class of regularized VAR models, called hierarchical vector autoregression (HVAR), that embed lag order selection into a convex regularizer to provide more reliable forecasts and more interpretable output.

The HVAR shifts the focus from obtaining estimates that are generally sparse (as measured by the number of nonzero autoregressive coefficients) to attaining estimates with *low maximal lag order*. While our motivating goal is to produce interpretable models with improved forecast performance, a convenient byproduct of the HVAR framework is a flexible and computationally efficient method for lag order selection.

Lag order selection procedures have been developed since the inception of VAR. Early attempts utilize least squares estimation with an information criterion or hypothesis testing (Lütkepohl 1985). The asymptotic theory of these approaches is well developed in the fixed-dimensional setting, in which the length of the series T grows while the number of components k and the maximal lag order p are held fixed (White 2001); however, for small sample sizes, it has been observed that no criterion works well (Nickelsburg 1985). Gonzalo & Pitarakis (2002) find that for fixed k and p , when T is relatively small, Akaike’s Information Criterion (AIC) tends to overfit whereas Schwartz’s Information Criterion (BIC) tends to severely underfit. Despite their shortcomings, AIC, BIC, and corrected AIC (AICc, Hurvich & Tsai 1989) are still the preferred tools for lag order selection by most practitioners (Lütkepohl 2007, Pfaff 2008, Tsay 2013).

A drawback with such approaches is that they typically require the strong assumption of a single, universal lag order that applies across all components. While this reduces the computational complexity of model selection, it has little statistical or economic justification, it unnecessarily constrains the dynamic relationship between the components, and it impedes forecast performance. Gredenhoff & Karlsson (1999) show that violation of the universal lag order assumption can lead to overparameterized models or the imposition of false zero restrictions. They instead suggest considering *componentwise* specifications that allow each marginal regression to have a different lag order (which is sometimes referred to as an *asymmetric VAR*). One such procedure (Hsiao 1981) starts from univariate autoregressive models and sequentially adds lagged components according to Akaike’s “Final Prediction Error” criterion (Akaike 1969). However, this requires an *a priori* ranking of components based on their perceived predictive power, which is inherently subjective. Keating (2001) offers a more general method which estimates all potential p^k componentwise VAR specifications and utilizes AIC or BIC for lag order selection. Such an approach is computationally intractable and standard asymptotic justifications are inapplicable if the number of components k is large. Ding & Karlsson (2014) present several specifications which allow for varying lag order within a Bayesian framework. Markov chain Monte Carlo estimation methods with spike and slab priors are proposed, but these are computationally intensive, and estimation becomes intractable in high dimensions.

Given the difficulties with lag order selection in VAR models, many authors have turned instead to shrinkage-based approaches, which impose sparsity, or other economically-motivated restrictions, on the parameter space to make reliable estimation tractable. Early shrinkage methods, such as Litterman (1979), take a pragmatic Bayesian perspective. Many such

approaches apply the *Minnesota prior*, which uses natural conjugate priors to shrink the VAR toward either an intercept-only model or toward a vector random walk, depending on the context. The prior covariance is specified so as to incorporate the belief that a series' *own* lags are more informative than *other* lags and that lower lags are more informative than higher lags. With this prior structure, coefficients with high lags will have a prior mean of zero and a prior variance that decays with the lag. Hence, coefficients with higher lags are shrunk more toward zero; however, as in ridge regression, coefficients will not be estimated as exactly zero.

More recent shrinkage approaches have incorporated the lasso (Tibshirani 1996). Hsu et al. (2008) considers the lasso with common information criterion methods for model selection. The use of the lasso mitigates the need to conduct an exhaustive search over the space of all $2^{k^2 p}$ possible models but does not explicitly encourage lags to be small. Song & Bickel (2011) use a group lasso (Yuan & Lin 2006) penalty structure and down-weight higher lags via scaling the penalty parameter by an increasing function of the coefficients' lag. The authors note that the functional form of these weights is arbitrary, but the estimates are sensitive to the choice of weights.

In Section 2 we introduce the HVAR framework, which attacks the traditional lag order selection problem through convex regularization. HVAR forces low lag coefficients to be selected before corresponding high lag coefficients, thereby specifically shrinking toward low lag order solutions. This is in contrast to approaches such as Song & Bickel (2011), which increase the weight of the penalty parameter with the coefficients' lag without explicitly enforcing a low-lag structure. In Section 2.1 we introduce three hierarchical lag structures that may be desirable when fitting VAR models to data. These structures vary in the degree to which lag order selection is common across different components. For each lag structure, a corresponding HVAR model is detailed in Section 2.2 for attaining that sparsity structure. The proposed methodology allows for flexible HVAR estimation in the high dimensional setting with a single tuning parameter. We develop algorithms in Section 3 that are computationally efficient and parallelizable across components. A simulation study in Section 4 and two macroeconomic applications in Section 5 highlight the advantages of HVAR in forecasting and lag order selection. The appendix provides additional details of our simulation methodology.

2 Methodology

Let $\{\mathbf{y}_t \in \mathbb{R}^k\}_{t=1}^T$ denote a k -dimensional vector time series of length T . A p th order vector autoregression $\text{VAR}_k(p)$ may be expressed as a multivariate regression

$$\mathbf{y}_t = \boldsymbol{\nu} + \boldsymbol{\Phi}^{(1)}\mathbf{y}_{t-1} + \cdots + \boldsymbol{\Phi}^{(p)}\mathbf{y}_{t-p} + \mathbf{u}_t, \quad \text{for } t = 1, \dots, T, \quad (2.1)$$

conditional on initial values $\{\mathbf{y}_{-(p-1)}, \dots, \mathbf{y}_0\}$, in which $\boldsymbol{\nu} \in \mathbb{R}^k$ denotes an intercept vector, $\{\boldsymbol{\Phi}^{(\ell)} \in \mathbb{R}^{k \times k}\}_{\ell=1}^p$ denote lag- ℓ coefficient matrices, and $\{\mathbf{u}_t \in \mathbb{R}^k\}_{t=1}^T$ denotes a mean zero white noise (serially uncorrelated) vector time series with an unspecified $k \times k$ nonsingular contemporaneous covariance matrix $\boldsymbol{\Sigma}_u$.

In the classical low-dimensional setting in which $T > kp$, one may perform least squares to fit the $\text{VAR}_k(p)$ model,

minimizing

$$\sum_{t=1}^T \|\mathbf{y}_t - \boldsymbol{\nu} - \sum_{\ell=1}^p \boldsymbol{\Phi}^{(\ell)} \mathbf{y}_{t-\ell}\|_2^2 \quad (2.2)$$

over $\boldsymbol{\nu}$ and $\{\boldsymbol{\Phi}^{(\ell)}\}$, where $\|\mathbf{a}\|_2 = (\sum_i \mathbf{a}_i^2)^{1/2}$ denotes the Euclidean norm of a vector \mathbf{a} . We will find it convenient to express the VAR using compact matrix notation:

$$\begin{aligned} \mathbf{Y} &= [\mathbf{y}_1 \cdots \mathbf{y}_T] & (k \times T); & \quad \boldsymbol{\Phi} = [\boldsymbol{\Phi}^{(1)} \cdots \boldsymbol{\Phi}^{(p)}] & (k \times kp); \\ \mathbf{z}_t &= [\mathbf{y}_{t-1}^\top \cdots \mathbf{y}_{t-p}^\top]^\top & (kp \times 1); & \quad \mathbf{Z} = [\mathbf{z}_1 \cdots \mathbf{z}_T] & (kp \times T); \\ \mathbf{U} &= [\mathbf{u}_1 \cdots \mathbf{u}_T] & (k \times T); & \quad \mathbf{1} = [1 \cdots 1]^\top & (T \times 1). \end{aligned}$$

Equation (2.1) is then simply

$$\mathbf{Y} = \boldsymbol{\nu} \mathbf{1}^\top + \boldsymbol{\Phi} \mathbf{Z} + \mathbf{U},$$

and the least squares procedure (2.2) can be expressed as minimizing

$$\|\mathbf{Y} - \boldsymbol{\nu} \mathbf{1}^\top - \boldsymbol{\Phi} \mathbf{Z}\|_2^2$$

over $\boldsymbol{\nu}$ and $\boldsymbol{\Phi}$, where $\|\mathbf{A}\|_2$ denotes the Frobenius norm of the matrix \mathbf{A} , that is the Euclidean norm of $\text{vec}(\mathbf{A})$ (not be mistaken for the operator norm, which does not appear in this paper).

Estimating the parameters of this model is challenging unless T is sufficiently large. We therefore seek to incorporate reasonable structural assumptions on the parameter space to make estimation tractable for moderate to small T . Multiple authors have considered using the lasso penalty, building in the assumption that the lagged coefficient matrices $\boldsymbol{\Phi}^{(\ell)}$ are sparse (Song & Bickel 2011, Davis et al. 2012, Hsu et al. 2008); theoretical work has elucidated how such structural assumptions can lead to better estimation performance even when the number of parameters is large (Basu & Michailidis 2013). In what follows, we define a class of sparsity patterns, which we call hierarchical lag structures, that arises in the context of multivariate time series.

2.1 Hierarchical Lag Structures

In Equation (2.1), the parameter $\boldsymbol{\Phi}_{ij}^{(\ell)}$ controls the dynamic dependence of the i th component of \mathbf{y}_t on the j th component of $\mathbf{y}_{t-\ell}$. In describing hierarchical lag structures, we will use the following notational convention: for $1 \leq \ell \leq p$, let

$$\begin{aligned} \boldsymbol{\Phi}^{(\ell:p)} &= [\boldsymbol{\Phi}^{(\ell)} \cdots \boldsymbol{\Phi}^{(p)}] \in \mathbb{R}^{k \times k(p-\ell+1)} \\ \boldsymbol{\Phi}_i^{(\ell:p)} &= [\boldsymbol{\Phi}_i^{(\ell)} \cdots \boldsymbol{\Phi}_i^{(p)}] \in \mathbb{R}^{1 \times k(p-\ell+1)} \\ \boldsymbol{\Phi}_{ij}^{(\ell:p)} &= [\boldsymbol{\Phi}_{ij}^{(\ell)} \cdots \boldsymbol{\Phi}_{ij}^{(p)}] \in \mathbb{R}^{1 \times (p-\ell+1)}. \end{aligned}$$

Consider the $k \times k$ matrix of *elementwise* coefficient lags \mathbf{L} defined by

$$\mathbf{L}_{ij} = \max\{\ell : \Phi_{ij}^{(\ell)} \neq 0\},$$

in which we define $\mathbf{L}_{ij} = 0$ if $\Phi_{ij}^{(\ell)} = 0$ for all $\ell = 1, \dots, p$. Therefore, each \mathbf{L}_{ij} denotes the maximal coefficient lag (maxlag) for component j in the regression model for component i . In particular, \mathbf{L}_{ij} is the smallest ℓ such that $\Phi_{ij}^{([\ell+1]:p)} = \mathbf{0}$. Note that the maxlag matrix \mathbf{L} is not symmetric, in general. There are numerous hierarchical lag structures that one can consider within the context of the $\text{VAR}_k(p)$ model. The simplest such structure is that $\mathbf{L}_{ij} = L$ for all i and j , meaning that there is a *universal* (U) maxlag that is shared by every pair of components. Expressed in terms of Equation (2.1), this would say that $\Phi^{([L+1]:p)} = \mathbf{0}$ and that $\Phi_{ij}^{(L)} \neq 0$ for all $1 \leq i, j \leq k$. While the methodology we introduce can be easily extended to this and many other potential hierarchical lag structures, in this paper we focus on the following three fundamental structures.

1. **Componentwise (C).** A componentwise hierarchical lag structure allows each of the k marginal equations from (2.1) to have its own maxlag, but all components within each equation must share the same maximal lag:

$$\mathbf{L}_{ij} = L_i \quad \forall j, \quad \text{for } i = 1, \dots, k.$$

Hence in Equation (2.1), this implies $\Phi_i^{([L_i+1]:p)} = \mathbf{0}$ and $\Phi_{ij}^{(L_i)} \neq 0$ for all i and j . This componentwise hierarchical lag structure is illustrated in Figure 1.

2. **Own-Other (O).** The own-other hierarchical lag structure is similar to the componentwise structure, but with an added within-lag hierarchy that imposes the mild assumption that a series' own lags ($i = j$) are more informative than other lags ($i \neq j$). Thus, diagonal elements are prioritized before off-diagonal elements within each lag, componentwise (i.e., row-wise). In particular,

$$\mathbf{L}_{ij} = L_i^{\text{other}} \quad \text{for } i \neq j \quad \text{and} \quad \mathbf{L}_{ii} \in \{L_i^{\text{other}}, L_i^{\text{other}} + 1\}, \quad \text{for } i = 1, \dots, k.$$

This hierarchical lag structure allows each component of \mathbf{y}_t to have longer range lagged self-dependence than lagged cross-dependencies. This own-other hierarchical lag structure is illustrated in Figure 2.


3. **Elementwise (E).** Finally, we consider a completely flexible hierarchical lag structure in which the elements of \mathbf{L} have no stipulated relationships. This elementwise hierarchical lag structure is illustrated in Figure 3.

In the next section, we introduce the proposed class of HVAR estimators aimed at estimating $\text{VAR}_k(p)$ models while shrinking the elements of \mathbf{L} towards zero by incorporating the three hierarchical lag structures described above.

2.2 HVAR: Hierarchical Group Lasso for Lag Structured VAR Modeling


In this section, we introduce convex penalties specifically tailored for attaining the three lag structures presented in the previous section. Our primary modeling tool is the hierarchical group lasso (Zhao et al. 2009), which is a group lasso (Yuan & Lin 2006) with a nested group structure. The group lasso is a sum of (unsquared) Euclidean norms and is used in statistical

Figure 1: A componentwise (C) hierarchical lag structure: $\text{HVAR}_3^C(5)$.




$$\mathbf{L}^C = \begin{pmatrix} 5 & 5 & 5 \\ 2 & 2 & 2 \\ 4 & 4 & 4 \end{pmatrix}$$

Figure 2: An own-other (O) hierarchical lag structure: $\text{HVAR}_3^O(5)$.



$$\mathbf{L}^O = \begin{pmatrix} 5 & 4 & 4 \\ 2 & 2 & 2 \\ 3 & 3 & 4 \end{pmatrix}$$

Figure 3: An elementwise (E) hierarchical lag structure: $\text{HVAR}_3^E(5)$.



$$\mathbf{L}^E = \begin{pmatrix} 5 & 3 & 4 \\ 5 & 0 & 5 \\ 5 & 5 & 2 \end{pmatrix}$$

modeling as a penalty to encourage groups of parameters to be set to zero simultaneously. Using nested groups leads to hierarchical sparsity constraints in which one set of parameters being zero implies that another set is also zero. This penalty has been applied to multiple statistical problems including regression models with interactions (Zhao et al. 2009, Jenatton et al. 2010, Radchenko & James 2010, Bach et al. 2012, Bien et al. 2013, Lim & Hastie 2013, Haris et al. 2014, She & Jiang 2014), covariance estimation (Bien et al. 2014), additive modeling (Lou et al. 2014), and time series (Suo & Tibshirani 2014). This last work focuses on transfer function estimation, in this case scalar regression with multiple time-lagged covariates whose coefficients decay with lag.

For each hierarchical lag structure presented above, we propose an estimator based on a convex optimization problem.

1. The \mathbf{HVAR}^C objective is a *componentwise* hierarchical lag structure and is defined by

$$\min_{\nu, \Phi} \left\{ \frac{1}{2} \|\mathbf{Y} - \nu \mathbf{1}^\top - \Phi \mathbf{Z}\|_2^2 + \lambda \sum_{i=1}^k \sum_{\ell=1}^p \|\Phi_i^{(\ell:p)}\|_2 \right\}, \quad (2.3)$$

in which $\|\mathbf{A}\|_2$ denotes the Euclidean norm of $\text{vec}(\mathbf{A})$, for a matrix \mathbf{A} . As the penalty parameter $\lambda \geq 0$ is increased, we have $\hat{\Phi}_i^{(\ell:p)} = \mathbf{0}$ for more i , and for smaller ℓ . This componentwise hierarchical group structure builds in the condition that if $\hat{\Phi}_i^{(\ell)} = 0$, then $\hat{\Phi}_i^{(\ell')} = 0$ for all $\ell' > \ell$, for each $i = 1, \dots, k$. This structure favors lower maxlag models componentwise, rather than simply giving sparse Φ estimates with no particular structure.

2. The \mathbf{HVAR}^O objective aims for a *own-other* hierarchical lag structure and is defined by

$$\min_{\nu, \Phi} \left\{ \frac{1}{2} \|\mathbf{Y} - \nu \mathbf{1}^\top - \Phi \mathbf{Z}\|_2^2 + \lambda \sum_{i=1}^k \sum_{\ell=1}^p \left[\|\Phi_i^{(\ell:p)}\|_2 + \|(\Phi_{i,-i}^{(\ell)}, \Phi_i^{([\ell+1]:p)})\|_2 \right] \right\}, \quad (2.4)$$

in which $\Phi_{i,-i}^{(\ell)} = \{\Phi_{ij}^{(\ell)} : j \neq i\}$, and where we adopt the convention that $\Phi_i^{([p+1]:p)} = \mathbf{0}$. The first term in this penalty is identical to that of (2.3). The difference is the addition of the second penalty term, which is just like the first except that it omits $\Phi_{ii}^{(\ell)}$. This penalty allows sparsity patterns in which the influence of component i on itself may be nonzero at lag ℓ even though the influence of other components is thought to be zero at that lag. This model ensures that, for all $\ell' > \ell$, $\hat{\Phi}_i^{(\ell)} = \mathbf{0}$ implies $\hat{\Phi}_i^{(\ell')} = \mathbf{0}$ and $\hat{\Phi}_{ii}^{(\ell)} = \mathbf{0}$ implies $\hat{\Phi}_{i,-i}^{(\ell'+1)} = \mathbf{0}$. This accomplishes the desired own-other hierarchical lag structure such that $\mathbf{L}_{i,-i} = L_i^{other} \mathbf{1}_{k-1}$ and $\mathbf{L}_{ii} \in \{L_i^{other}, L_i^{other} + 1\}$, componentwise.

3. The **HVAR**^E objective aims for an *elementwise* hierarchical lag structure and is defined by

$$\min_{\nu, \Phi} \left\{ \frac{1}{2} \|\mathbf{Y} - \nu \mathbf{1}^\top - \Phi \mathbf{Z}\|_2^2 + \lambda \sum_{i=1}^k \sum_{j=1}^k \sum_{\ell=1}^p \|\Phi_{ij}^{(\ell:p)}\|_2 \right\}. \quad (2.5)$$

Here, each of the k^2 pairs of components can have its own maxlag, such that $\Phi_{ij}^{(\ell:p)} = \mathbf{0}$ may occur for different values of ℓ for each pair i and j . While this model is the most flexible of the three, it also borrows the least strength across the different components. When \mathbf{L}_{ij} differ for all i and j , we expect this method to do well, whereas when, for example $\mathbf{L}_{ij} = L_i$, we expect it to be inefficient relative to (2.3).

Since all three objectives are based on hierarchical group lasso penalties, a unified computational approach to solve each is detailed in the next section.

3 Optimization Algorithm

We begin by noting that since the intercept ν does not appear in the penalty terms, it can be removed if we replace \mathbf{Y} by $\mathbf{Y}(\mathbf{I}_T - \frac{1}{T} \mathbf{1} \mathbf{1}^\top)$ and \mathbf{Z} by $\mathbf{Z}(\mathbf{I}_T - \frac{1}{T} \mathbf{1} \mathbf{1}^\top)$. All three optimization problems are of the form

$$\min_{\Phi} \left\{ \frac{1}{2} \|\mathbf{Y} - \Phi \mathbf{Z}\|_2^2 + \lambda \sum_{i=1}^k \sum_{\ell=1}^p \Omega_i(\Phi_i^{(\ell:p)}) \right\}, \quad (3.1)$$

and (2.3), (2.4), and (2.5) only differ by the form of the norm Ω_i . A key simplification is possible by observing that the objective above decouples across the rows of Φ :

$$\min_{\Phi} \sum_{i=1}^k \left[\frac{1}{2} \|\mathbf{Y}_i - \Phi_i \mathbf{Z}\|_2^2 + \lambda \sum_{\ell=1}^p \Omega_i(\Phi_i^{(\ell:p)}) \right],$$

in which $\mathbf{Y}_i \in \mathbb{R}^{1 \times T}$ and $\Phi_i = \Phi_i^{(1:p)} \in \mathbb{R}^{1 \times kp}$. Hence, Equation (3.1) can be solved in parallel by solving the “one-row” subproblem

$$\min_{\Phi_i} \left\{ \frac{1}{2} \|\mathbf{Y}_i - \Phi_i \mathbf{Z}\|_2^2 + \lambda \sum_{\ell=1}^p \Omega_i(\Phi_i^{(\ell:p)}) \right\}. \quad (3.2)$$

Jenatton et al. (2011) show that hierarchical group lasso problems can be efficiently solved via the proximal gradient method. This procedure can be viewed as an extension of traditional gradient descent methods to nonsmooth objective functions.

Given a convex objective function of the form $f_i(\Phi_i) = \mathcal{L}_i(\Phi_i) + \lambda\Omega_i^*(\Phi_i)$, where \mathcal{L}_i is differentiable with a Lipschitz continuous gradient, the proximal gradient method produces a sequence $\hat{\Phi}_i[1], \hat{\Phi}_i[2], \dots$ with the guarantee that

$$f_i(\hat{\Phi}_i[m]) - \min_{\Phi_i} f_i(\Phi_i)$$

is $O(1/m)$ (cf. Beck & Teboulle 2009). For $m = 1, 2, \dots$, its update is given by

$$\hat{\Phi}_i[m] = \text{Prox}_{s_m \lambda \Omega_i^*} \left(\hat{\Phi}_i[m-1] - s_m \nabla \mathcal{L}_i(\hat{\Phi}_i[m-1]) \right),$$

where s_m is an appropriately chosen step size and $\text{Prox}_{s_m \lambda \Omega_i^*}$ is the proximal operator of the function $s_m \lambda \Omega_i^*(\cdot)$, which is evaluated at the gradient step we would take if we were minimizing \mathcal{L}_i alone. The proximal operator is defined as the unique solution of a convex optimization problem involving Ω_i^* but not \mathcal{L}_i :

$$\text{Prox}_{s_m \lambda \Omega_i^*}(u) = \underset{v}{\operatorname{argmin}} \left\{ \frac{1}{2} \|u - v\|_2^2 + s_m \lambda \Omega_i^*(v) \right\}. \quad (3.3)$$

The proximal gradient method is particularly effective when the proximal operator can be evaluated efficiently. In our case, $\Omega_i^*(\Phi_i) = \sum_{\ell=1}^p \Omega_i(\Phi_i^{(\ell;p)})$ is a sum of hierarchically nested Euclidean norms. Jenatton et al. (2011) show that for such penalties, the proximal operator has essentially a closed form solution, making it extremely efficient. It remains to note that $\mathcal{L}_i(\Phi_i) = \frac{1}{2} \|\mathbf{Y}_i - \Phi_i \mathbf{Z}\|_2^2$ has gradient $\nabla \mathcal{L}_i(\Phi_i) = -(\mathbf{Y}_i - \Phi_i \mathbf{Z}) \mathbf{Z}^\top$ and that the step size s_m can be determined adaptively through a backtracking procedure or it can be set to the Lipschitz constant of $\nabla \mathcal{L}_i(\Phi_i)$, which in this case is $\sigma_1(\mathbf{Z})^{-2}$ (where $\sigma_1(\mathbf{Z})$ denotes the largest singular value of \mathbf{Z}).

Beck & Teboulle (2009) develop an accelerated version of the proximal gradient method, which they call the Fast Iterative Soft-Thresholding Algorithm (FISTA). This leads to a faster convergence rate and improved empirical performance with minimal additional overhead. Our particular implementation is based on Algorithm 2 of Tseng (2008). It repeats, for $m = 1, 2, \dots$ to convergence,

$$\begin{aligned} \hat{\phi} &\leftarrow \hat{\Phi}_i[m-1] + \frac{m-2}{m+1} \left(\hat{\Phi}_i[m-1] - \hat{\Phi}_i[m-2] \right) \\ \hat{\Phi}_i[m] &\leftarrow \text{Prox}_{s_m \lambda \Omega_i^*} \left(\hat{\phi} - s_m \nabla \mathcal{L}_i(\hat{\phi}) \right), \end{aligned}$$

and converges at rate $1/m^2$ (compared to the unaccelerated proximal gradient method's $1/m$ rate). The full procedure is detailed in Algorithm 1 and is applicable to all three HVAR estimators.

The algorithms for these methods differ only in the evaluation of their proximal operators (since each method has a different penalty Ω_i^*). However, all three choices of Ω_i^* correspond to hierarchical group lasso penalties, allowing us to use the result of Jenatton et al. (2011), which shows that the proximal operator has a remarkably simple form. We write these three problems generically as

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \left\{ \frac{1}{2} \|\mathbf{x} - \tilde{\mathbf{x}}\|_2^2 + \lambda \sum_{h=1}^H w_h \|\mathbf{x}_{g_h}\|_2 \right\}, \quad (3.4)$$

Algorithm 1 General algorithm for HVAR with penalty Ω_i^*

Require: $\mathbf{Y}, \mathbf{Z}, \hat{\Phi}[0], \lambda, \epsilon$
 $\hat{\Phi}[1] \leftarrow \hat{\Phi}[0]; \quad \hat{\Phi}[2] \leftarrow \hat{\Phi}[0]$
 $s \leftarrow \sigma_1(\mathbf{Z})^{-2}$
for $i = 1, \dots, k$ **do**
 for $m = 1, 2, \dots$ **do**
 $\hat{\phi} \leftarrow \hat{\Phi}_i[m-1] + \frac{m-2}{m+1} (\hat{\Phi}_i[m-1] - \hat{\Phi}_i[m-2])$
 $\hat{\Phi}_i[m] \leftarrow \text{Prox}_{s\lambda\Omega_i^*} \left(\hat{\phi} + s \cdot (\mathbf{Y}_i - \hat{\phi}\mathbf{Z})\mathbf{Z}^\top \right)$
 if $\|\hat{\phi} - \hat{\Phi}_i[m]\|_\infty \leq \epsilon$ **then**
 break
 end if
 end for
end for
return $\hat{\Phi}[m]$

where $g_1 \subset \dots \subset g_H$. The key observation in Jenatton et al. (2011) is that the dual of the proximal problem (3.3) can be solved exactly in a *single pass* of blockwise coordinate descent. By strong duality, this solution to the dual provides us with a solution to problem (3.3). Furthermore, the updates of each block are extremely simple, corresponding to a groupwise-soft-thresholding operation. Algorithm 2 shows the solution to (3.4), which includes all three of our penalties as special cases.

Algorithm 2 Solving Problem (3.4)

Require: $\tilde{\mathbf{x}}, \lambda, w_1, \dots, w_H$
 $\mathbf{r} \leftarrow \tilde{\mathbf{x}}$
for $h = 1, \dots, H$ **do**
 $\mathbf{r}_{g_h} \leftarrow (1 - \lambda w_h / \|\mathbf{r}_{g_h}\|_2) \mathbf{r}_{g_h}$
end for
return \mathbf{r} as the solution $\hat{\mathbf{x}}$.

4 Simulation Study

In this section we compare the proposed HVAR methods with competing VAR modeling approaches. After detailing these comparison methods below, we describe three simulation scenarios and discuss the forecast and lag order selection performance of each estimator. Finally, we examine the performance of the proposed HVAR methods in a low-dimensional simulation setting while allowing the maximal lag order p to vary.

4.1 Comparison Methods

A standard method in lower dimensional settings is to fit a $\text{VAR}_k(\ell)$ with least squares for $0 \leq \ell \leq p$ and then to select a universal lag order ℓ using AIC or BIC. Per Lütkepohl (2007), the AIC and BIC of a $\text{VAR}_k(\ell)$ are defined as

$$\text{AIC}(\ell) = \log \det(\hat{\Sigma}_u^\ell) + \frac{2k^2\ell}{T}, \quad (4.1)$$

$$\text{BIC}(\ell) = \log \det(\hat{\Sigma}_u^\ell) + \frac{\log(T)k^2\ell}{T}, \quad (4.2)$$

in which $\hat{\Sigma}_u^\ell$ is the residual sample covariance matrix having used least squares to fit the $VAR_k(\ell)$. The lag order ℓ that minimizes $AIC(\ell)$ or $BIC(\ell)$ is selected. This method of lag order selection is only possible when $k\ell \leq T$ since otherwise least squares is not well-defined. In the first set of simulations that follow, we cannot use least squares for $\ell > 1$, thus for a simple benchmark we instead estimate a $VAR_k(1)$ by least squares:

$$\min_{\nu, \Phi} \left\{ \frac{1}{2} \|\mathbf{Y} - \nu \mathbf{1}^\top - \Phi^{(1)} \mathbf{Z}^{(1)}\|_2^2 \right\},$$

where $\mathbf{Z}^{(1)} = [\mathbf{y}_0 \cdots \mathbf{y}_{T-1}]$. We also include two well-known regularization approaches. The *lasso* estimates the VAR using an L_1 -penalty:

$$\min_{\nu, \Phi} \left\{ \frac{1}{2} \|\mathbf{Y} - \nu \mathbf{1}^\top - \Phi \mathbf{Z}\|_2^2 + \lambda \|\Phi\|_1 \right\},$$

where $\|\Phi\|_1$ denotes $\|\text{vec}(\Phi)\|_1$. The lasso does not intrinsically consider lag order, hence Song & Bickel (2011) propose a *lag-weighted lasso* penalty in which a weighted L_1 -penalty is used with weights that increase geometrically with lag order:

$$\min_{\nu, \Phi} \left\{ \frac{1}{2} \|\mathbf{Y} - \nu \mathbf{1}^\top - \Phi \mathbf{Z}\|_2^2 + \lambda \sum_{\ell=1}^p \ell^\alpha \|\Phi^{(\ell)}\|_1 \right\}.$$

The tuning parameter $\alpha \in [0, 1]$ determines how fast the penalty weight increases with lag. While this form of penalty applies greater regularization to higher order lags, it is less structured than our HVAR penalties in that it does not necessarily produce sparsity patterns in which all coefficients beyond a certain lag order are zero.

Finally, we compare against two naive approaches to serve as simple baselines: the unconditional *sample mean* corresponds to the intercept-only model,

$$\min_{\nu} \frac{1}{2} \|\mathbf{Y} - \nu \mathbf{1}^\top\|_2^2,$$

which makes one-step-ahead forecasts of the form $\hat{\mathbf{y}}_{t+1} = \frac{1}{t} \sum_{\ell=1}^t \mathbf{y}_\ell$; and the vector *random walk* model, which corresponds to

$$\hat{\nu} = \mathbf{0}, \quad \hat{\Phi}^{(1)} = \mathbf{I}_k, \quad \hat{\Phi}^{(2:p)} = \mathbf{0},$$

and makes one-step-ahead forecasts of the form $\hat{\mathbf{y}}_{t+1} = \mathbf{y}_t$.

4.2 Simulation Settings

In order to demonstrate the efficacy of the HVAR methods in applications with various lag structures, we evaluate forecasts produced by the proposed methods under several simulation scenarios. In these scenarios, we have $k = 60$ components, a maximal lag order of $p = 12$, and a series length of $T = 100$; the error covariance is taken to be $\Sigma_u = 0.01 \cdot \mathbf{I}_{60}$. All simulations are generated from stationary coefficient matrices. The steps taken to ensure the stationarity of the simulation structures are described in detail in Section A.1 of the appendix. Simulation results are based on 100 iterations.

The penalty parameters were selected using the rolling cross-validation approach utilized by Song & Bickel (2011) and Banbura et al. (2009), with the middle third of the data used for penalty parameter selection and the last third for forecast evaluation. In the case of the lag-weighted lasso, λ and α were jointly selected. Given an evaluation period $(T_1, T_2]$, we use

Figure 4: *Sparsity pattern (and magnitudes) of the $HVAR_{60}^C(5)$ structure used in simulation Scenario 1.*




Table 1: *Out-of-sample mean-squared one-step-ahead forecast error (standard errors are in parentheses) for Scenario 1 based on 100 simulations.*

Class	Method	MSFE
HVAR	Componentwise	0.0464 (0.0005)
	Own-other	0.0470 (0.0005)
	Elementwise	0.0546 (0.0005)
VAR	Lasso	0.0598 (0.0006)
	Lag-weighted lasso	0.0547 (0.0006)
	Least squares	0.1599 (0.0020)
Other	Sample mean	0.1295 (0.0034)
	Random walk	0.3054 (0.0107)

mean-squared one-step-ahead forecast error (MSFE) as a measure of forecast performance:

$$MSFE(T_1, T_2) = \frac{1}{k(T_2 - T_1)} \sum_{i=1}^k \sum_{t=T_1}^{T_2-1} (\hat{\mathbf{y}}_{i,t+1} - \mathbf{y}_{i,t+1})^2,$$

where $\hat{\mathbf{y}}_{i,t+1}$ denotes the forecast of a method for time $t + 1$ and component i based on observing the series up to time t .

We evaluate the methods under three lag structures.

Simulation Scenario 1: Componentwise Lag Structure. In this scenario, we simulate according to an $HVAR_{60}^C(5)$ structure. In particular, we choose the maxlag matrix

$$\mathbf{L} = [1, 2, 3, 4, 5]^\top \otimes (\mathbf{1}_{12} \mathbf{1}_{60}^\top).$$

This 60×60 maxlag matrix is row-wise constant, meaning that all components *within a row* have the same maxlag; we partition the rows into 5 groups of size 12, each group taking on a distinct maxlag in $\{1, 2, 3, 4, 5\}$. A coefficient matrix Φ with maxlag matrix \mathbf{L} is used in Scenario 1’s simulations and its magnitudes are depicted in Figure 8. The prediction performance of the methods under study is shown in Table 1. The componentwise and own-other HVAR methods perform best, which is to be expected since both methods are geared explicitly toward a lag structure as in Scenario 1. The lag-weighted lasso and elementwise HVAR perform similarly, and both of them do better than the regular lasso. With a total of $pk^2 = 12(60)^2 = 43,200$ coefficient parameters to estimate, the methods that assume an ordering are greatly advantaged over a method like the lasso that does not exploit this knowledge. One exception is the $VAR_{60}(1)$ model that is fit using least squares: Despite this method’s explicit orientation toward modeling recent behavior, it suffers both because it misses important longer range lag coefficients and because it is an unregularized estimator of $\Phi^{(1)}$ and therefore has high variance.

Simulation Scenario 2: Own-Other Lag Structure. In this scenario, we create the matrix Φ in such a manner that it differentiates between *own* and *other* coefficients. The coefficients of a series’ “own lags” (i.e., $\Phi_{ii}^{(\ell)}$) are larger in magnitude

Figure 5: *Sparsity pattern (and magnitudes) of the $\text{HVAR}_{60}^O(2)$ structure used in simulation Scenario 2.*




Table 2: *Out-of-sample mean-squared one-step-ahead forecast error (standard errors are in parentheses) for Scenario 2 based on 100 simulations.*

Class	Method	MSFE
HVAR	Componentwise	0.0193 (0.0001)
	Own-other	0.0183 (0.0001)
	Elementwise	0.0210 (0.0002)
VAR	Lasso	0.0270 (0.0003)
	Lag-weighted lasso	0.0228 (0.0003)
	Least squares	0.0544 (0.0007)
Other	Sample mean	0.2948 (0.0178)
	Random walk	0.1381 (0.0086)

than those of “other lags” (i.e., $\Phi_{ij}^{(\ell)}$ with $i \neq j$). The magnitude of coefficients decreases as the lag order increases. The $\text{HVAR}_{60}^O(2)$ model we simulate is depicted in Figure 5. The first 20 rows can be viewed as univariate autoregressive models in which only the *own* term is nonzero; in the next 20 rows, for the first k coefficients, the coefficient on a series’ *own* lags is larger than “other lags,” and, for the next k coefficients, only *own* coefficients are nonzero; the final 20 rows have nonzeros throughout the first $2k$ coefficients, with *own* coefficients dominating *other* coefficients in magnitude. The results from this scenario are shown in Table 2. Here, all three HVAR methods lead the competing methods. As one would expect, the own-other HVAR procedure achieves the best forecasting performance, with the componentwise HVAR performing only slightly worse. We find that the lag-weighted lasso performs worse than the elementwise HVAR, but much better than the lasso without weights. As with the previous scenario, the least-squares approach is not competitive.

Simulation Scenario 3: Elementwise Lag Structure. In this final scenario, we simulate under an $\text{HVAR}_{60}^E(4)$ model, meaning that the maxlag is allowed to vary not just across rows but also *within* rows. The maxlag matrix is given by

$$\mathbf{L} = \begin{pmatrix} 4 & 3 & 2 & 1 \\ 3 & 2 & 1 & 0 \\ 2 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} \otimes (\mathbf{1}_{15} \mathbf{1}_{15}^\top).$$

A coefficient matrix corresponding to this lag structure is depicted in Figure 6, and the results are shown in Table 3.

As expected, the elementwise HVAR method outperforms all others. The chosen \mathbf{L} violates the own-other lag structure in 45 of the 60 rows (and it violates the componentwise lag structure in every row). Even so, these two HVAR methods outperform the lasso and weighted lasso methods.

Figure 6: *Sparsity pattern (and magnitudes) of the $HVAR_{60}^E(4)$ structure used in simulation Scenario 3.*




Table 3: *Out-of-sample mean-squared one-step-ahead forecast error (standard errors are in parentheses) for Scenario 3 based on 100 simulations.*

Class	Method	MSFE
HVAR	Componentwise	0.0323 (0.0007)
	Own-other	0.0335 (0.0007)
	Elementwise	0.0295 (0.0004)
VAR	Lasso	0.0379 (0.0007)
	Lag-weighted lasso	0.0380 (0.0007)
	Least squares	0.1071 (0.0016)
Other	Sample mean	0.0915 (0.0038)
	Random walk	0.1253 (0.0039)

4.3 Lag Order Selection

While our primary intent in introducing the HVAR framework is better forecast performance and improved interpretability, one can also view HVAR as an approach for selecting lag order. Below, we examine the performance of the proposed methods in estimating the maxlag matrix \mathbf{L} defined in Section 2.1. Based on an estimate $\hat{\Phi}$ of the autoregressive coefficients, we can likewise define a matrix of estimated lag orders:

$$\hat{\mathbf{L}}_{ij} = \max\{\ell : \hat{\Phi}_{ij}^{(\ell)} \neq 0\},$$

where we define $\hat{\mathbf{L}}_{ij} = 0$ if $\hat{\Phi}_{ij}^{(\ell)} = 0$ for all ℓ . It is well known in the regularized regression literature (cf., Leng et al. 2006) that the optimal tuning parameter for prediction is different from that for support recovery. Nonetheless, in this section we will proceed with the rolling cross-validation procedure used previously with only two minor modifications intended to ameliorate the tendency of cross-validation to select a value of λ that is smaller than optimal for support recovery. First, we cross-validate a relaxed version of the regularized methods in which the estimated nonzero coefficients are refit using ridge regression. This refitting procedure is described in detail in Section A.2 in the appendix. This modification makes the MSFE more sensitive to $\hat{\mathbf{L}}_{ij}$ being larger than necessary. Second, we use the “one-standard-error rule” heuristic discussed in Hastie et al. (2009), in which we select the largest value of λ whose MSFE is no more than one standard error above that of the best performing model (since we favor the most parsimonious model that does approximately as well as any other).

We measure a procedure’s lag order selection performance based on the sum of absolute differences between \mathbf{L} and $\hat{\mathbf{L}}$:

$$\|\hat{\mathbf{L}} - \mathbf{L}\|_1 = \sum_{ij} |\hat{\mathbf{L}}_{ij} - \mathbf{L}_{ij}|.$$

In particular, Tables 4, 5, and 6 report this value for various methods relative to that of the sample mean (which chooses

Table 4: *Lag selection performance (standard errors in parentheses) for Scenario 1 based on 100 simulations.*

Class	Method	$\ \hat{\mathbf{L}} - \mathbf{L}\ _1 / \ \mathbf{L}\ _1$
HVAR	Componentwise	0.5898 (0.0495)
	Own-other	0.5792 (0.0447)
	Elementwise	0.9684 (0.0034)
VAR	Lasso	0.9968 (0.0013)
	Lag-weighted lasso	0.9936 (0.0023)
Other	Sample mean	1.0000 (0.0000)

Table 5: *Lag selection performance (standard errors in parentheses) for Scenario 2 based on 100 simulations.*

Class	Method	$\ \hat{\mathbf{L}} - \mathbf{L}\ _1 / \ \mathbf{L}\ _1$
HVAR	Componentwise	0.5376 (0.0060)
	Own-other	0.4329 (0.0061)
	Elementwise	0.9561 (0.0009)
VAR	Lasso	1.0395 (0.0023)
	Lag-weighted lasso	1.0551 (0.0031)
Other	Sample mean	1.0000 (0.0000)

$\hat{\mathbf{L}}_{ij} = 0$ for all i and j).

In Scenario 1, the own-other and componentwise HVARs achieve the best performance; every other approach scarcely outperforms the benchmark. The fact that the own-other and componentwise HVARs perform best is no surprise given that they both are geared toward a lag structure as in Scenario 1.

In Scenario 2, the own-other HVAR achieves the best performance followed by the componentwise HVAR; the elementwise HVAR performs much worse than these but still better than all other methods.

Interestingly, in Scenario 3, the elementwise HVAR approach is the only method to perform better than the sample-mean baseline. We see that the HVAR methods that incorrectly assume that maxlag should be constant (or near constant) within a row pay a price in lag order selection, making them even worse than the lasso methods.

4.4 Simulation Scenario 4: Robustness of HVAR as p increases

We additionally examine the impact of the upper bound for maximal lag order p on HVAR's performance. Ideally, provided that p is large enough to capture the dynamics of the system, its choice should have little impact on forecast performance. However, we should expect regularizers that treat each coefficient democratically, such as the lasso, to experience degraded forecast performance as p increases.

As an experiment, we simulate from an $\text{HVAR}_{10}^C(5)$ while increasing the upper bound on the maximal lag order to substantially exceed the true \mathbf{L} . All series in the first 4 rows have $\mathbf{L} = 2$, the next 3 rows have $\mathbf{L} = 5$, and the final 3 rows

Table 6: *Lag selection performance (standard errors in parentheses) for Scenario 3 based on 100 simulations.*

Class	Method	$\ \hat{\mathbf{L}} - \mathbf{L}\ _1 / \ \mathbf{L}\ _1$
HVAR	Componentwise	1.3260 (0.0382)
	Own-other	1.2573 (0.0365)
	Elementwise	0.8782 (0.0035)
VAR	Lasso	1.0121 (0.0020)
	Lag-weighted lasso	1.0144 (0.0043)
Other	Sample mean	1.0000 (0.0000)

Figure 7: Sparsity pattern (and magnitudes) of the $HVAR_{10}^C(5)$ structure used in simulation Scenario 4.





Figure 8: Simulation Results: Scenario 4 (AIC omitted due to extremely poor performance)



have $\mathbf{L} = \mathbf{0}$. Figure 7 depicts the coefficient matrix Φ and its magnitudes.

We consider varying $p \in (1, 5, 12, 25, 50)$. As p increases, we should expect the performance of the HVAR procedures to remain relatively constant whereas the lasso and information-criterion based methods should return worse forecasts.

At $p = 1$ all models are misspecified. Since no method is capable of capturing the true dynamics of series 1-7 in Figure 8, all perform poorly. As expected, after ignoring $p = 1$, the componentwise HVAR achieves the best performance across all other choices for p , although own-other and elementwise HVAR's performances are within one standard error. Among the information-criterion based methods, AIC performs substantially worse than BIC as p increases. This is likely the result of BIC assigning a larger penalty on the number of coefficients than AIC. The lasso's performance degrades substantially as the lag order increases, while the lag-weighted lasso is somewhat more robust to the lag order, but still achieves worse forecasts than every HVAR procedure under all scenarios except for $p = 25$ where it performs comparably.

Table 7: Out-of-sample mean-squared one-step-ahead forecast error (standard errors in parentheses) for Scenario 4 based on 100 simulations ($T=200$).

Class	Method	MSFE (p=1)	MSFE (p=5)	MSFE (p=12)	MSFE (p=25)	MSFE (p=50)
HVAR	Componentwise	0.0141 (0.0010)	0.0119 (0.0077)	0.0120 (0.0008)	0.0122 (0.0009)	0.0125 (0.0010)
	Own-other	0.0142 (0.0010)	0.0120 (0.0008)	0.0121 (0.0008)	0.0123 (0.0008)	0.0126 (0.0010)
	Elementwise	0.0142 (0.0011)	0.0120 (0.0008)	0.0121 (0.0008)	0.0123 (0.0008)	0.0126 (0.0009)
VAR	Lasso	0.0142 (0.0011)	0.0125 (0.0008)	0.0138 (0.0010)	0.0156 (0.0013)	0.0190 (0.0016)
	Lag-weighted lasso	0.0143 (0.0011)	0.0121 (0.0008)	0.0122 (0.0008)	0.0123 (0.0015)	0.0129 (0.0012)
	AIC	0.0141 (0.0010)	0.0117 (0.0080)	0.0535 (0.0070)	0.0781 (0.0121)	0.0855 (0.0130)
	BIC	0.0141 (0.0010)	0.0140 (0.0011)	0.0141 (0.0011)	0.0142 (0.0012)	0.0144 (0.0013)

5 Data Analysis

5.1 Macroeconomic Application

We now apply the proposed HVAR estimation methods to a collection of US macroeconomic time series compiled by Stock & Watson (2005) and augmented by Koop (2011). The full dataset contains 168 quarterly macroeconomic indicators over 45 years, representing information about many aspects of the US economy, including income, employment, stock prices, exchange rates, etc. The full list of variables considered is available in Koop (2011), who defines various nested groups of components:

- *Small* ($k = 3$): Federal Funds Rate, CPI, and GDP growth rate; this core group is commonly considered in basic Dynamic Stochastic Generalized Equilibrium modeling;
- *Medium* ($k = 20$): Small group plus 17 additional variables, including indices for consumption, labor, and housing, as well as exchange rates;
- *Medium-Large* ($k = 40$): Medium group plus 20 additional aggregate variables;
- *Large* ($k = 168$): Medium-Large group plus 128 additional variables, consisting primarily of the components that make up the aggregated variables (e.g. subsets of Gross Domestic Product, Bond Interest Rates, Industrial Production, etc).

Forecast Comparisons. We initially focus on forecasting the *Medium-Large* ($k = 40$) and *Large* ($k = 168$) groups. We apply the transformation codes provided by Stock & Watson (2005) to make the data approximately stationary, then we standardize each series to have sample mean zero and variance one. Quarter 3 of 1977 to Quarter 3 of 1992 is used for penalty parameter selection while Quarter 4 of 1992 to Quarter 4 of 2007 is used for forecast performance comparisons. Following the convention from Banbura et al. (2009), we set the maximal lag order p to 13. In the *Large* group, VAR by AIC and BIC are overparameterized and not included.

The rolling out of sample one-step-ahead mean square forecast error (MSFE) for the *Medium-Large* and *Large* groups are summarized in Table 9. The proposed HVAR methods outperformed the lasso, least squares, and both information-criterion based models for the *Medium-Large* group over this evaluation period. Among the HVAR methods, the more flexible own-other and elementwise structures performed similarly, and better than the componentwise structure. The sample mean and random walk forecast results are provided for comparison.

As the number of component series k increases, the componentwise hierarchical lag order structure becomes less realistic. This is especially true in high-dimensional economic applications, in which a core subset of the included series is typically most important in forecasting. In Table 9 we see that the componentwise HVAR performs more similarly to the lasso and least squares methods for the *Large* group over this evaluation period. The own-other and elementwise HVAR methods again had the best forecasting performance. This supports the widely held belief that in economic applications, a components' *own* lags are likely more informative than *other* lags and that maxlag varies across components.

Lag Order Selection. The *Small* group includes Real Gross Domestic Product (GDP251), a measure of economic activity, Consumer Price Index (CPIAUSL), a measure of inflation, and the Federal Funds Rate (FYFF), a measure of monetary policy. This core subset is generally of primary interest to forecasters and policymakers. We now examine the estimated lag

Table 8: *Rolling out of sample one-step ahead MSFE for the Medium-Large ($k = 40$) and Large ($k = 168$) groups of macroeconomic indicators.*

Class	Method	Medium-Large $k = 40$	Large $k = 168$
HVAR	Componentwise	0.5342	0.6188
	Own-other	0.5102	0.5749
	Elementwise	0.5138	0.5752
VAR	Lasso	0.5385	0.5916
	Lag-weighted lasso	0.5262	0.5876
	AIC	2.9750	N/A
	BIC	0.7502	N/A
Other	Sample mean	0.6861	0.7486
	Random walk	1.1775	1.3306

order of these three component series from a fitted HVAR_{40}^E model of the *Medium-Large* group over the entire observation period. The estimated lag order is shown in Figure 11.

Most of the lag orders chosen by the elementwise HVAR have an underlying economic interpretation. The Federal Funds Rate (FYFF), has been shown in Bernanke & Blinder (1992) to be an important predictor of several measures of economic activity, including the components of Gross Domestic Product. Additionally, the “Taylor Rule” (Taylor 1993) suggests that the Federal Funds Rate is set to control inflation, hence we should expect changes in the previous quarter to aid in forecasting inflation.

The non-core component series which have high maximal lag orders are also economically rational. *FYGT10*, the interest rate on 10 year maturity Treasury Bonds, has a maxlag of 6 in the Consumer Price Index regression and 3 in the Federal Funds Rate regression. 10 year bond yields historically serve as a proxy for the Federal Reserve’s monetary policy, and here we see that it aids in predicting both inflation and the federal funds rate. *UTL11*, industrial capacity utilization, is an important indicator of economic activity, and similar to inflation, it is widely believed that the Federal Reserve sets its Federal Funds Rate in order to achieve a target level of capacity utilization (McElhattan 1978). *GDP263*, which denotes real exports has a maxlag of 8 in the GDP growth rate regression; Marin (1992) showed that there exists a Granger causal relationship between the US growth rate and net exports. *GDP273* is a price index constructed by the Bureau of Economic Analysis which has objectives similar to the Consumer Price Index, hence, they appear to exhibit a high degree of lagged dependence.

5.2 Exchange Rate Application

We additionally consider utilizing the proposed HVAR procedures to forecast monthly exchange rates. Exchange rates are notoriously difficult to forecast; most economically-motivated models are substantially outperformed by a simple random walk (Meese & Rogoff 1983). However, recent developments by Carriero et al. (2009) have demonstrated that jointly forecasting a panel of exchange rates using a Bayesian VAR can lead to substantial forecast improvements over a random walk.


Our application forecasts the monthly exchange rates of the currencies of four South American nations in relation to the US dollar. These currencies include the Argentinian Peso (ARS), Brazilian Real (BRL), Chilean Peso (CLP), and the Peruvian Nuevo Sol (PEN). We chose a cross-section of emerging currencies rather than their more established counterparts as emerging currencies are subjected to a wide range of monetary policies but are still closely related, both economically and

Table 9: *Rolling out of sample one-step ahead MSFE for $k = 4$ monthly exchange rate forecasts (relative to a random walk), $p = 12$*

Class	Method	Relative MSFE
HVAR	Componentwise	0.885
	Own-other	0.878
	Elementwise	0.914
VAR	Lasso	0.935
	Lag-weighted lasso	0.943
	AIC	1.064
	BIC	1.093
Other	Sample mean	44.00
	Random walk	1.000

geographically. Our data were acquired from the proprietary CUR database hosted on **Quandl** and ranges from June of 2003 to January 2015 ($T = 140$). The series are plotted in Figure 9.

Figure 9: *Plots of the monthly exchange rate vis-a-vis the US dollar for the Brazilian Real (BRL, first), the Peruvian Nuevo Sol (PEN, second), Argentinian Peso (ARS, third), and the Chilean Peso (CLP, fourth).*



Note that the Argentinian Peso and Peruvian Nuevo Sol both follow strikingly different trends than the other two currencies. This is likely due to their respective nations' aggressive inflation-targeting monetary policies which require intense intervention to promote stability as compared to the relatively free-floating policies of the other two countries, in which rates tend to vary according to market demand (Frenkel & Rapetti 2010).

Since all of these series exhibit evidence of nonstationarity, instead of shrinking every coefficient to zero, following Carriero et al. (2009), we modify our procedures to instead shrink toward a vector random walk. The training period ranges from November 2006 to June 2010 and the forecast evaluation ranges from July 2010 to January 2015. Table 9 displays the rolling out of sample one-step-ahead MSFE (relative to a random walk).


We find that every HVAR and lasso VAR model improves upon the random walk's forecasting performance, with the own-other HVAR and componentwise HVAR achieving the greatest improvement. These results suggest that even at a low dimension, imposing the HVAR framework leads to a substantial improvement in forecasting performance.

As in the previous application, we examine the estimated lag orders selected by the elementwise HVAR for each of the

exchange rates, which are plotted in Figure 10. Though the true underlying exchange rate dynamics are likely beyond the scope of this model, the maxlags selected offer interesting economic insights. There is a strong bilateral relationship between Argentina and Peru (maxlags of 6 and 5, respectively). This could be the result of a strong degree of interdependence between their economies, as both economies are centered around exports of raw materials. Argentina’s exchange rate appears to be very important in forecasting Brazil’s (maxlag of 12), but the reverse is not true. This unilateral relationship could be due to Brazil’s considerable imports from Argentina.

Additionally, there appears to be a strong bilateral relationship between Brazil and Chile (maxlags of 9 and 4, respectively). Though Brazil officially practices a “managed floating” exchange rate, in practice its government intervenes heavily to control the exchange rate in an effort to protect its international competitiveness (Leahy 2012). It is possible that the importance of Chile’s exchange rates is the result of Brazil’s interventionist response to global macroeconomic conditions that affect both economies.

Figure 10: *Plot of \hat{L}^E , denoting the estimated elementwise maxlag for each exchange rate series.*



6 Discussion

By incorporating the property that more recent lags convey more information than distant lags, the hierarchical VAR approaches offer substantial forecast improvements as well as greater insight into lag order selection than existing methods. In addition, throughout our simulation scenarios, we see that each method is fairly robust to deviations from its particular hierarchical structure. The substantial improvements in forecasting accuracy in data applications provide justification for the widely held belief that as the number of component series included in a model increases, the maximal lag order is not symmetric across series. Our methods scale well and are computationally feasible in high dimensions. Implementations of our methods are available in the R package **BigVAR**, which is hosted on Github.

Figure 11: The first three rows of \hat{L}^E , denoting the estimated elementwise maxlag for each series in the Medium-Large group using the $HVAR^E$ method. Components with maxlag of zero are left empty. The first component, Federal Funds Rate (FYFF), has been shown in Bernanke & Blinder (1992) to be an important predictor of several measures of economic activity, including the components of Gross Domestic Product. Additionally, the “Taylor Rule” (Taylor 1993) suggests that the Federal Funds Rate is set to control inflation, hence we should expect changes in the previous quarter to aid in forecasting inflation.

	GDP251	CPIAUCSL	FYFF	UTL11	GDP273	FM1	FYGT10	PMDEL	GDP263
GDP251	2		2	1		2	1		
CPIAUCSL	3	1	1	1	1	2	8	2	1
FYFF	1	1	1	5	1	6	1	3	1

References

- Akaike, H. (1969), ‘Fitting autoregressive models for prediction’, *Annals of the institute of Statistical Mathematics* **21**(1), 243–247.
- Bach, F., Jenatton, R., Mairal, J. & Obozinski, G. (2012), ‘Structured sparsity through convex optimization’, *Statistical Science* **27**(4), 450–468.
- Back, A. D. & Weigend, A. S. (1997), ‘A First Application of Independent Component Analysis to Extracting Structure from Stock Returns’, *International Journal of Neural Systems* **8**, 473–484.
- Banbura, M., Giannone, D. & Reichlin, L. (2009), ‘Large bayesian vector auto regressions’, *Journal of Applied Econometrics* **25**(1), 71–92.
- Basu, S. & Michailidis, G. (2013), ‘Estimation in high-dimensional vector autoregressive models’, *arXiv preprint arXiv:1311.4175*.
- Beck, A. & Teboulle, M. (2009), ‘A fast iterative shrinkage-thresholding algorithm for linear inverse problems’, *SIAM Journal on Imaging Sciences* **2**(1), 183–202.
- Bernanke, B. S. & Blinder, A. S. (1992), ‘The federal funds rate and the channels of monetary transmission’, *The American Economic Review* pp. 901–921.
- Bien, J., Bunea, F. & Xiao, L. (2014), ‘Convex banding of the covariance matrix’, *arXiv preprint arXiv:1405.6210*.
- Bien, J., Taylor, J. & Tibshirani, R. (2013), ‘A lasso for hierarchical interactions’, *The Annals of Statistics* **41**(3), 1111–1141.
- Boshnakov, G. N. & Iqelan, B. M. (2009), ‘Generation of time series models with given spectral properties’, *Journal of Time Series Analysis* **30**(3), 349–368.
- Box, G. E. P. & Tiao, G. C. (1977), ‘A Canonical Analysis of Multiple Time Series’, *Biometrika* **64**(2), 355.
- Carriero, A., Kapetanios, G. & Marcellino, M. (2009), ‘Forecasting exchange rates with a large bayesian var’, *International Journal of Forecasting* **25**(2), 400–417.
- Davis, R. A., Zang, P. & Zheng, T. (2012), ‘Sparse vector autoregressive modeling’. journal: arXiv preprint arXiv:1207.0520.
- Ding, S. & Karlsson, S. (2014), Bayesian var models with asymmetric lags, Technical report, Orebro University, Orebro University School of Business, Orebro University, Sweden.
- Forni, M., Hallin, M., Lippi, M. & Reichlin, L. (2000), ‘The Generalized Dynamic-Factor Model: Identification and Estimation’, *Review of Economics and Statistics* **82**(4), 540–554.
- Frenkel, R. & Rapetti, M. (2010), ‘A concise history of exchange rate regimes in latin america’.
- Gilbert, P. (2005), ‘Brief users guide: Dynamic systems estimation (dse)’, *Available in the file doc/dse-guide. pdf distributed together with the R bundle dse, to be downloaded from <http://cran.r-project.org>*.

- Gonzalo, J. & Pitarakis, J.-Y. (2002), ‘Lag length estimation in large dimensional systems’, *Journal of Time Series Analysis* **23**(4), 401–423.
- Gredenhoff, M. & Karlsson, S. (1999), ‘Lag-length selection in var-models using equal and unequal lag-length procedures’, *Computational Statistics* **14**(2), 171–187.
- Haris, A., Witten, D. & Simon, N. (2014), ‘Convex modeling of interactions with strong heredity’, *arXiv preprint arXiv:1410.3517*.
- Hastie, T., Tibshirani, R., Friedman, J., Hastie, T., Friedman, J. & Tibshirani, R. (2009), *The elements of statistical learning*, Springer.
- Hsiao, C. (1981), ‘Autoregressive modelling and money-income causality detection’, *Journal of Monetary economics* **7**(1), 85–106.
- Hsu, N. J., Hung, H. L. & Chang, Y. M. (2008), ‘Subset selection for vector autoregressive processes using lasso’, **52**(7), 3645–3657. journal: Computational Statistics & Data Analysis.
- Hurvich, C. M. & Tsai, C.-L. (1989), ‘Regression and time series model selection in small samples’, *Biometrika* **76**(2), 297–307.
- Jenatton, R., Mairal, J., Obozinski, G. & Bach, F. (2010), Proximal methods for sparse hierarchical dictionary learning, in ‘Proceedings of the International Conference on Machine Learning (ICML)’.
- Jenatton, R., Mairal, J., Obozinski, G. & Bach, F. (2011), ‘Proximal methods for hierarchical sparse coding’, *The Journal of Machine Learning Research* **12**, 2297–2334.
- Keating, J. W. (2001), ‘Macroeconomic modeling with asymmetric vector autoregressions’, *Journal of Macroeconomics* **22**(1), 1–28.
- Koop, G. (2011), ‘Forecasting with medium and large bayesian vars’, *Journal of Applied Econometrics*.
- Leahy, J. (2012), ‘Brazil admits tight hold over exchange rate’, *Financial Times*.
- Leng, C., Lin, Y. & Wahba, G. (2006), ‘A note on the lasso and related procedures in model selection’, *Statistica Sinica* **16**(4), 1273.
- Lim, M. & Hastie, T. (2013), ‘Learning interactions through hierarchical group-lasso regularization’, *arXiv preprint arXiv:1308.2719*.
- Litterman, R. B. (1979), Techniques of forecasting using vector autoregressions, Working Papers 115, Federal Reserve Bank of Minneapolis.
URL: <http://ideas.repec.org/p/fip/fedmwpl/115.html>
- Lou, Y., Bien, J., Caruana, R. & Gehrke, J. (2014), ‘Sparse partially linear additive models’, *arXiv preprint arXiv:1407.4729*.

- Lütkepohl, H. (1985), ‘Comparison of criteria for estimating the order of a vector autoregressive process’, *Journal of time series analysis* **6**(1), 35–52.
- Lütkepohl, H. (2007), *New introduction to multiple time series analysis*, Springer.
- Marin, D. (1992), ‘Is the export-led growth hypothesis valid for industrialized countries?’, *The Review of Economics and Statistics* pp. 678–688.
- Matteson, D. S. & Tsay, R. S. (2011), ‘Dynamic orthogonal components for multivariate time series’, *Journal of the American Statistical Association* **106**(496), 1450–1463.
URL: <http://pubs.amstat.org/doi/abs/10.1198/jasa.2011.tm10616>
- McElhattan, R. (1978), ‘Estimating a stable-inflation capacity-utilization rate’, *Economic Review* **Fall**, 20–30.
- Meese, R. A. & Rogoff, K. (1983), ‘Empirical exchange rate models of the seventies: Do they fit out of sample?’, *Journal of international economics* **14**(1), 3–24.
- Meinshausen, N. (2007), ‘Relaxed lasso’, *Computational Statistics & Data Analysis* **52**(1), 374–393.
- Neumaier, A. & Schneider, T. (2001), ‘Estimation of parameters and eigenmodes of multivariate autoregressive models’, *ACM Transactions on Mathematical Software (TOMS)* **27**(1), 27–57.
- Nickelsburg, G. (1985), ‘Small-sample properties of dimensionality statistics for fitting var models to aggregate economic data: A monte carlo study’, *Journal of Econometrics* **28**(2), 183–192.
- Peña, D. & Box, G. E. P. (1987), ‘Identifying a Simplifying Structure in Time Series’, *Journal of the American Statistical Association* **82**(399), 836–843.
- Pfaff, B. (2008), *Analysis of integrated and cointegrated time series with R*, Springer.
- Radchenko, P. & James, G. M. (2010), ‘Variable selection using adaptive nonlinear interaction structures in high dimensions’, *Journal of the American Statistical Association* **105**(492), 1541–1553.
- She, Y. & Jiang, H. (2014), ‘Group Regularized Estimation under Structural Hierarchy’, *ArXiv e-prints* .
- Sims, C. A. (1980), ‘Macroeconomics and reality’, *Econometrica: Journal of the Econometric Society* pp. 1–48.
- Song, S. & Bickel, P. (2011), ‘Large vector auto regressions’. journal: arXiv preprint arXiv:1106.3915.
- Stock, J. H. & Watson, M. W. (2002), ‘Forecasting Using Principal Components from a Large Number of Predictors’, *Journal of the American Statistical Association* **97**(460), 1167–1179.
- Stock, J. H. & Watson, M. W. (2005), ‘An empirical comparison of methods for forecasting using many predictors’, *Manuscript, Princeton University* .
- Suo, X. & Tibshirani, R. (2014), ‘An ordered lasso and sparse time-lagged regression’, *arXiv preprint arXiv:1405.6447* .

- Taylor, J. B. (1993), Discretion versus policy rules in practice, *in* ‘Carnegie-Rochester conference series on public policy’, Vol. 39, Elsevier, pp. 195–214.
- Tiao, G. C. & Tsay, R. S. (1989), ‘Model Specification in Multivariate Time Series’, *Journal of the Royal Statistical Society. Series B (Methodological)* **51**(2), 157–213.
- Tibshirani, R. (1996), ‘Regression shrinkage and selection via the lasso’, *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 267–288.
- Tsay, R. S. (2013), *Multivariate Time Series Analysis: With R and Financial Applications*, John Wiley & Sons.
- Tseng, P. (2008), ‘On accelerated proximal gradient methods for convex-concave optimization’, *submitted to SIAM Journal on Optimization*.
- White, H. (2001), *Asymptotic theory for econometricians*, Academic press New York.
- Yuan, M. & Lin, Y. (2006), ‘Model selection and estimation in regression with grouped variables’, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **68**(1), 49–67.
- Zhao, P., Rocha, G. & Yu, B. (2009), ‘The composite absolute penalties family for grouped and hierarchical variable selection’, *The Annals of Statistics* **37**(6A), 3468–3497.

A Appendix

A.1 Generation of Simulation Scenarios

All of our simulation structures were generated to ensure a stationary coefficient matrix, Φ . In order to construct a coefficient matrix for these scenarios, we started by converting the $\text{VAR}_k(p)$ to a $\text{VAR}_k(1)$ as described in equation 2.1.8 of Lütkepohl (2007)

$$\mathbf{A} = \begin{bmatrix} \Phi^{(1)} & \Phi^{(2)} & \dots & \Phi^{(p-1)} & \Phi^{(p)} \\ \mathbf{I}_k & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_k & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I}_k & \mathbf{0} \end{bmatrix} \quad (\text{A.1})$$

For \mathbf{A} to be stationary, its maximum eigenvalue must be less than 1. In general, it is very difficult to generate stationary coefficient matrices. Boshnakov & Iqelan (2009) offers a potentially viable procedure that utilizes the unique structure of A.1, but it does not allow for structured sparsity. We instead follow the approach put forth by Gilbert (2005) in which structured random coefficient matrices are generated until a stationary matrix is recovered.

A.2 Relaxed VAR Estimation

Since the Lasso and its structured counterparts are known to shrink non-zero regression coefficients, in practice, they are often used for model selection, followed by refitting the reduced model using least squares (Meinshausen 2007). In this section, we detail our approach to refit based on the support selected by our procedures while taking into consideration both numerical stability as well as computational efficiency.

Let $\hat{\mathbf{B}}$ denote the coefficient matrix recovered from one of our algorithms and suppose that it contains r nonzero coefficients. In order to take the support recovered into account we introduce \mathbf{V} , a $k^2p \times r$ *restriction matrix* of rank r that denotes the location of nonzero elements in $\hat{\Phi}$. Defining β as the vec of the nonzero entries of $\hat{\Phi}$, we obtain the relationship

$$\text{vec}(\hat{\Phi}) = \mathbf{V}\beta.$$

We can then express the *Relaxed Least Squares* estimator as:

$$\text{vec}(\hat{\Phi}_{\text{Relaxed}}) = \mathbf{V}[\mathbf{V}^\top(\mathbf{Z}\mathbf{Z}^\top \otimes \mathbf{I}_k)\mathbf{V}]^{-1}\mathbf{V}^\top(\mathbf{Z} \otimes \mathbf{I}_k)\text{vec}(\mathbf{Y}), \quad (\text{A.2})$$

in which \otimes denotes the Kronecker operator. In general, it is ill-advised to directly form A.2. First, performing matrix operations with $\mathbf{Z} \otimes \mathbf{I}_k$, which has dimension $kT \times k^2p$, can be very computationally demanding, especially if k is large. Second, in the event that $r \approx T$, the resulting estimator can be very poorly conditioned. To obviate these two concerns, we propose a slight adaptation of the techniques detailed in Neumaier & Schneider (2001) that computes a variant of Equation A.2 using a QR decomposition to avoid explicit matrix inversion. Additionally, if the resulting matrix is found to be ill-conditioned, a small ridge penalty should be utilized to ensure numerically-stable solutions.

A.3 Refinements

As opposed to performing a Kronecker expansion we instead consider imposing the restrictions by row in $\hat{\Phi}$ and define V_1, \dots, V_k as $kp \times r_i$ restriction matrices of rank r_1, \dots, r_k , denoting the number of nonzero elements in each row of Φ . We can then calculate each row of $\hat{\Phi}_{\text{Relaxed}}$ by

$$\hat{\Phi}_{\text{Relaxed}_i} = (V_i(V_i^\top \mathbf{Z}\mathbf{Z}^\top V_i)^{-1}V_i^\top \mathbf{Z}\mathbf{Y}_i)^\top. \quad (\text{A.3})$$

Now, following Neumaier & Schneider (2001), construct the matrix $\mathbf{K}_i = [(V_i\mathbf{Z})^\top, \mathbf{Y}_i]$. We then compute a QR factorization of \mathbf{K}_i

$$\mathbf{K}_i = \mathbf{Q}\mathbf{R},$$

in which Q is an orthogonal matrix and R is upper triangular of the form:

$$R = \begin{bmatrix} r_i & 1 \\ R_{11} & R_{12} \\ 0 & R_{22} \end{bmatrix} \begin{matrix} \\ r_i \\ T-r_i \end{matrix}$$

As expanded upon in Neumaier & Schneider (2001), we can compute

$$\begin{aligned} \hat{\Phi}_{\text{Relaxed}_i} &= (V_i R_{12}^\top R_{11} (R_{11}^\top R_{11})^{-1})^\top, \\ &= (V_i R_{12}^\top R_{11} R_{11}^{-1} (R_{11}^\top)^{-1})^\top, \\ &= (V_i R_{12}^\top (R_{11}^\top)^{-1})^\top, \\ &= (V_i (R_{11}^{-1} R_{12})^\top)^\top, \end{aligned}$$

which can be evaluated with a triangular solver, hence does not require explicit matrix inversion. In the event that \mathbf{K} is poorly conditioned, to improve numerical stability, we add a small ridge penalty. It is suggested by Neumaier & Schneider (2001) to add a penalty corresponding to scaling a diagonal matrix D consisting of the Euclidean norms of the columns of \mathbf{K} by $(r_i^2 + r_i + 1)\epsilon_{\text{machine}}$. The full refitting algorithm is detailed in Algorithm 3.

Algorithm 3 Relaxed Least Squares

Require: $\mathbf{Z}, \mathbf{Y}, V_1, \dots, V_k$
for $i = 1, 2, \dots, k$ **do**
 $\mathbf{K}_i \leftarrow [(V_i \mathbf{Z})^\top, \mathbf{Y}_i]$
 $D \leftarrow (r_i^2 + r_i + 1)\epsilon_{\text{machine}} \text{diag}(\|\mathbf{K}_i\|_2)$
 $R, Q \leftarrow QR\left(\begin{bmatrix} \mathbf{K}_i \\ D \end{bmatrix}\right)$
 $\hat{\Phi}_{\text{Relaxed}_i} \leftarrow (V_i (R_{11}^{-1} R_{12})^\top)^\top$
end for
return $\hat{\Phi}_{\text{Relaxed}}$.
